How to design institutions that foster cooperation

Arunas L. Radzvilavicius^{1,*}, Taylor A. Kessinger^{1,*}, Joshua B. Plotkin^{1,†}

¹Department of Biology, University of Pennsylvania, Philadelphia, PA, USA

* Equal contribution [†]E-mail: jplotkin@sas.upenn.edu

November 6, 2020

Abstract

Humans typically consider altruism a moral good and condition their social behavior on the moral reputations of others. Indirect reciprocity explains how social norms and moral reputations can collectively support cooperation: members of the society cooperate who others who are considered good. But the theory of indirect reciprocity requires institutions that monitor and publicly broadcast moral reputations. Here we study evolution of adherence to public monitoring in societies where individuals are, at first, independently responsible for evaluating the moral reputations of their peers. We show that adherence to a public institution of moral assessment can evolve and promote cooperation under all simple social norms, depending upon the institution’s tolerance to occasional antisocial behavior. We specify how an institution’s size and its degree of tolerance towards antisocial behavior can be designed to dramatically increase cooperation rates, even for social norms that previous studies have found to perform poorly. Public monitoring serves to eliminate disagreements about reputations in the population, which in turn increases cooperation and individual payoffs, so that adherence to the public institution can evolve by social contagion. Adherence is then robust to invasion, and it tends to crowd out internal mechanisms that could alternatively support cooperation. These results help explain why societies tend to elect centralized institutions to provide top-down moral governance of their individual behavior.

Introduction

Social norms, reputations, and institutions of moral assessment are critical to cooperation in human societies [Alexander, 1985, Elster, 1989, Kandori, 1992, Tomasello and Vaish, 2013, Hechter, 2018, Curry et al., 2019]. People typically consider altruistic behavior a moral good [Rand et al., 2013], and they tend to be more cooperative when being observed by others [Rege and Telle, 2004, Bereczkei et al., 2007]. Moral reputations are directly related to cooperation, and cooperation in turn can lead to higher social status [von Rueden et al., 2019], while kindness towards people of bad moral standing is sometimes punished [Martin et al., 2019]. Interactions in modern societies often involve cooperation with strangers, and so people must rely on moral reputation scores provided by institutions or third parties, for instance, in e-commerce interactions [Utz et al., 2009, Hechter, 2018]. Historically, communities have constructed institutions that broadcast information about others’ prior behavior so that all reputations are publicly available [Milgrom et al., 1990, Greif, 2004].

Evolutionary game theory provides a convenient framework for studying human behavior governed by social norms, reputations, and community enforcement [Nowak and Sigmund, 2005]. In such models of indirect reciprocity, a donor’s action (to cooperate or not) depends on the recipient’s moral reputation (“good” or “bad”). Reputations are updated according to a social

norm – a collection of rules that prescribe how an individual’s reputation depends on her past behavior towards others [Ohtsuki and Iwasa, 2004, Santos et al., 2018]. A simple social norm called Stern Judging, for example, assigns a good reputation to those who refuse to help individuals of bad moral standing, and to those who cooperate with other good members of the society [Kandori, 1992, Pacheco et al., 2006].

According to the large literature on indirect reciprocity, cooperation can evolve and remain stable under simple social norms like Stern Judging, as individuals adopt discriminatory behavior and cooperate only with players of good reputation [Nowak and Sigmund, 2005, Pacheco et al., 2006, Ohtsuki and Iwasa, 2006, Santos et al., 2016]. However, this explanation for cooperation relies on a public reputation system or rapid gossip [Nowak and Sigmund, 2005, Sommerfeld et al., 2007, Nakamura and Masuda, 2011, Ohtsuki et al., 2009], so there are no disagreements about the reputations of people among their peers. But in many realistic settings, people make their own subjective moral judgements about one another, and their views may differ due to different observation histories or independent errors in observing a focal individual’s actions [Smith and Apicella, 2019]. Under a framework of private reputations, even when everyone follows the same norm of moral assessment, cooperation tends to collapse due to disagreements in the population about each others’ moral standings [Okada et al., 2017, Hilbe et al., 2018].

One way to restore cooperation under private moral assessment is by empathetic perspective taking [Radzvilavicius et al., 2019, Masuda and Santos, 2019]. Although the reputation of a focal individual may vary according to the perspective of different observers (Figure 1a), empathy can reduce the rate of misunderstandings and unjustified defection. Moreover, empathy itself can evolve through social contagion, inducing high rates of cooperation typical of societies that enjoy an established public monitoring system. Nonetheless, empathetic perspective taking is not always be a feasible solution to the problem of cooperation under private moral evaluations. Evaluating social interactions from the perspective of another person carries a high cognitive cost, and not everyone is capable of empathizing or perspective-taking [Frith and Happé, 1999]. Moreover, inferences about the perspective of another person are not always accurate. And the benefits of empathy are vulnerable to the possibility of deception and manipulation [Smith, 2006]. For instance, a potential donor wishing to avoid the cost of cooperation with a recipient without being assigned a bad reputation could falsely portray a bad subjective assessment of all of her potential partners.

In order to avoid the cognitive costs and potential errors associated with empathetic moral evaluation, we hypothesize that a society might instead delegate moral evaluations to an institutional observer, or observers, who broadcast public and objective reputation scores for all members of the society. We define a moral institution as a sub-population of Q individuals, chosen to publicly provide their consensus view on the reputations of everyone in the population (Figure 1b). An institution can consist of a single individual, $Q = 1$, the whole society, $Q = N$, or any intermediate number of observers. The institution’s consensus view of a focal individual is determined by the mean reputation from the perspectives of the observers within the institution, along with a fixed strictness threshold q . In particular, the consensus public reputation of a focal individual is “good” provided a proportion q or more of the observers who form the institution see the focal individual as good. The threshold q is thus a measure of institutional strictness towards anti-social behavior. Under a strict institution (high q) even occasional antisocial behavior results in being assigned a bad public reputation, whereas under a tolerant institution (low q) occasional antisocial behavior is permitted without assigning a bad reputation.

Human societies have developed a variety of cultural institutions that reinforce conditional cooperation [Ostrom, 2000]. Here, we extend the theory of indirect reciprocity to study the evolutionary transition from individual, private moral assessment to institutionalized, public moral evaluation. We first investigate how adherence to different types of institutions for moral evalu-

Figure 1: Reputation matrices showing good (yellow) and bad (black) reputations of individuals (horizontal axis) as seen from the perspectives of different observers (vertical axis). (a) When moral evaluations are private, observers may disagree on an individual’s reputation. Disagreements arise from independent errors in observation or because different observers assess different social interactions involving a focal individual. (b) Instead of performing moral evaluations independently, people may instead choose to rely on reputations provided by a centralized institution, which in this figure consists of $Q = 2$ distinguished observers (first two rows) who broadcast their consensus views. When the entire population relies on the public institution for moral assessment, reputational disagreements are eliminated and cooperation can be facilitated.

ation determines equilibrium rates of cooperation, as behavioral strategies are allowed to evolve. Next we consider the evolution of adherence itself to a public institution of assessment. We allow individuals either to rely on reputations scores from their own subjective moral assessments of others, or alternatively to trust the reputation scores provided by the public institution. Depending upon the degree to which the institution tolerates antisocial behavior, we find that adherence to the public institution can spread by social contagion for all simple social norms and thereafter remain robust to invasion by individuals who ignore the institution. We conclude that adherence to public institutions of moral assessment can emerge naturally and that the rational design of institutions can achieve high rates of cooperation under all simple social norms. Finally, we discuss how public institutions may be formed by mutual agreement and what level of taxation a society will accept to support the operation of an institution of moral judgement.

Model

We study pairwise donation games with a payoff matrix of the form $\begin{pmatrix} b-c & -c \\ b & 0 \end{pmatrix}$. In any given interaction, if the donor chooses to cooperate with the recipient, she pays a cost c and her interaction partner receives a benefit $b > c$. However, if the donor defects, she pays no cost and the recipient receives no benefit.

Individuals make decisions based on their behavioral strategy. We consider three strategies: ALLC (always cooperate), ALLD (always defect), or DISC (cooperate only with recipients of good reputation). The frequencies of these three strategies in the population evolve via social contagion – that is, individuals tend to copy the strategies of more successful players (see Methods Section 1). Reputations are updated according to a social norm: a set of rules defining what types of behavior, observed by a third party, result in the donor being assigned a good or bad reputation [Sigmund,

2012, Ohtsuki and Iwasa, 2004]. While these rules can in general be complex, we focus on second-order normative rules that depend only on the donor’s action and the recipient’s reputation, since these norms tend to outperform more complex norms [Santos et al., 2018].

Specifically, we study the four social norms most prominent in the literature on indirect reciprocity: Stern Judging ($\begin{smallmatrix} G & B \\ B & G \end{smallmatrix}$), Simple Standing ($\begin{smallmatrix} G & B \\ G & G \end{smallmatrix}$), Scoring ($\begin{smallmatrix} B & B \\ B & G \end{smallmatrix}$), and Shunning ($\begin{smallmatrix} B & B \\ B & B \end{smallmatrix}$). The row index i indicates the donor’s action, $i = 1$ for defect or $i = 2$ for cooperate; and the column index j indicates the reputation of the recipient, $j = 1$ for bad or $j = 2$ for good. All four of these norms endorse (with a good reputation) cooperation with morally good recipients and condemn (with a bad reputation) defection with morally good recipients, but they differ in how they treat interactions with morally bad recipients. Simple Standing awards a good reputation for any interaction with a bad recipient, whereas Shunning awards a bad reputation for any interaction with a bad recipient. Stern Judging rewards defection with bad recipients and cooperation with good recipients. Finally, Scoring rewards cooperation with bad recipients and punishes defection with bad recipients, so that reputational assessments depend only on the action of the donor, not on the reputation of the recipient.

Our model also includes errors in strategy execution and in observation [Sasaki et al., 2017]: an individual who intends to cooperate can accidentally defect with probability e_1 , and an observer may erroneously assign a bad reputation instead of a good reputation, and *vice versa*, with probability e_2 .

We consider societies of individuals who are initially responsible for their own private moral evaluations of observed interactions: their evaluation of the donor depends on the donor’s action they observe, the social norm, and the recipient’s reputation. These evaluations are characterized by a level of empathy $0 \leq E \leq 1$: with probability $1 - E$, an observer will evaluate the donor using her own view of the recipient’s reputation, but with probability E , the observer will empathize with the donor, using instead the donor’s view of the recipient’s reputation [Radzvilavicius et al., 2019]. Alternatively, individuals can make their behavioral decisions using the public reputation of the potential recipient, as provided by a monitoring institution. The monitoring institution consists of $1 \leq Q \leq N$ distinguished observers who broadcast their consensus view of each person’s reputation. The public reputation of an individual is broadcast as “good” only if more than qQ institutional observers view her as good, and “bad” otherwise. Here $0 \leq q \leq 1$ is the strictness threshold of the institution, and it quantifies the institution’s degree of tolerance towards antisocial behavior: a lower value of q corresponds to a more permissive institution. First we study the effect of different institutions on levels of cooperation, and then we turn to the problem of whether adherence to a public institution will spread in a population.

Results

Cooperation under institutionalized moral assessment

We first investigated how different types of institutions for moral assessment influence cooperative behavior, assuming complete adherence to the public broadcast. The broad consensus in the literature on indirect reciprocity is that Stern Judging is the most socially beneficial norm of moral assessment, as it promotes the most cooperation among the second-order norms [Pacheco et al., 2006, Santos et al., 2016, 2018, Radzvilavicius et al., 2019]. In the hierarchy of norms, Stern Judging is closely followed by Simple Standing, whereas Shunning and Scoring are generally incapable of supporting cooperation rates exceeding 50% [Sasaki et al., 2017, Santos et al., 2016]. These prior results on social norms and cooperation assume a standard model of public information about reputations, which corresponds to a special case of our model: an institution with a single member

$Q = 1$.

In contrast to prior studies, we find that, regardless of the social norm adopted by the society, a public institution of moral assessment can be designed to support high levels of cooperation – as high or even higher than the level achieved under Stern Judging.

To establish these results, we use replicator dynamics [Taylor and Jonker, 1978] to describe changes over time in the frequencies of the three strategies (ALLC, ALLD and DISC) in an infinite population. Strategies spread or decline according to their payoffs relative to the population mean (see Methods Section 1). We assume that reputation updates reach equilibrium before individuals copy strategies of their more successful peers. That is, the timescales of reputation dynamics are faster than those of strategic change [Sasaki et al., 2017]. We contrast the strategic dynamics in societies that follow institutions to those in which individuals all form private moral judgments [Hilbe et al., 2018, Radzvilavicius et al., 2019]. Initially, we consider two extreme types of institutions: very strict institutions, which broadcast an individual as good only if all members agree she is good ($q > (Q - 1)/Q$), and very tolerant institutions, which broadcast an individual as good provided at least one member views her as good ($q < 1/Q$).

For all four second-order social norms considered, we find that public institutions of moral assessment generally support two strategic equilibria: a monomorphic population of unconditional defectors or a population consisting solely of discriminators (Figure 2). To quantify the likelihood that cooperation will evolve in such scenarios, we investigate the basins of attraction of cooperative equilibria, that is, the set of initial strategic states that leads to the dominance of discriminators and corresponding high levels of cooperation.

The basin of attraction towards the cooperative equilibrium depends on the social norm and on the institution’s degree of tolerance towards of antisocial behavior (Figure 2). Under the Scoring and Shunning norms, there are large basins of attraction toward cooperation when public reputations are provisioned by highly tolerant institutions, whereas neither strict institutional assessment nor private assessment support DISC as a stable equilibrium at all. Under Stern Judging and Simple Standing, both strict and stern public institutions support substantial basins of attraction towards cooperation, but the basins are even larger for strict institutions. Stern Judging, unlike Simple Standing, does not support any stable cooperation under (non-empathetic) private assessment.

Supplementary Figures 10, 11, and 12 demonstrate how basins of attraction towards cooperative play depend upon model parameters: the error rates e_1 and e_2 , as well as benefit-to-cost ratio b/c . Varying errors rates such either $e_1 \gg e_2$ or $e_1 \ll e_2$ has nearly imperceptible effects on the basins of attraction, for all social norms and types of institutions. By contrast, increasing the benefit-to-cost ratio b/c by a factor of two has a noticeable quantitative (but not qualitative) effect: it increases the size of any non-empty basin of attraction towards cooperation.

Analysis of replicator dynamics can help us understand cooperation rates in finite populations, as well. The evolution of strategies in small societies is additionally subject to demographic stochasticity and random strategy exploration at a rate μ . To study cooperation in this setting, we performed Monte Carlo simulations in which successful strategies spread via social contagion. That is, an individual copies the strategy with a probability $1 / (1 + \exp(-w[\Pi_1 - \Pi_0]))$, where w is the selection strength, and Π_1 and Π_0 are payoffs of two randomly selected individuals [Traulsen et al., 2007, 2010]. As in Radzvilavicius et al. [2019], we assumed that each individual plays many rounds of the donation game, with different partners, before reputations are updated and strategy imitation takes place (see Methods Section 3).

Figure 2: Public institutions of moral evaluation can facilitate the evolution of cooperation. We analyzed the evolution of behaviors in the donation game for societies with unempathetic private moral assessment (left column) compared to societies with two types of institutionalized moral judgment: tolerant ($q < \frac{1}{2}$, middle column) or strict ($q > \frac{1}{2}$, right column). The triangular simplex represents the frequencies of three behavioral strategies in an infinite population: ALLC (always cooperate), ALLD (always defect), and DISC (cooperate with good individuals, defect against bad), for each of four different social norms (rows). Blue arrows indicate the direction of strategic evolution under replicator dynamics, with circles indicating stable (filled) and unstable (empty) equilibria. The basins of attraction towards stable equilibria that support cooperation are highlighted in light blue. Error rates are $e_1 = e_2 = 0.02$, benefits and costs are $b = 5$, $c = 1$, and institution size is $Q = 2$. Analogous results for $e_1 \ll e_2$, $e_1 \gg e_2$, and $b = 10$ are shown in Supplementary Figures 10, 11, and 12.

Figure 3: Public institutions for moral assessment govern the level of cooperation as behavioral strategies are allowed to evolve. For each social norm, it is possible to specify a public institution of appropriate size (Q) and strictness (q) that induces far higher levels of stationary mean cooperation than in societies with private moral assessment. Public institutions can outperform even empathetic private assessment. The public institution that maximizes cooperation always requires more than $Q = 1$ individuals. Strict institutions (high q) perform best under Stern Judging and Simple Standing, whereas tolerant institutions (low q) perform best under Scoring and Shunning. Population size is $N = 50$, mutation rate is $\mu = 0.025$, and error rates are $e_1 = e_2 = 0.02$. The benefit and cost of cooperation are set at $b = 5$ and $c = 1$.

Monte Carlo simulations reveal that public institutions of moral assessment can support high levels of cooperation in finite populations (Figure 3). As suggested by our analysis of replicator dynamics (Figure 2), the stationary mean rate of cooperation in a finite population depends on both the institution size Q and its tolerance threshold q towards antisocial behavior. Importantly, for every social norm we study, we find that an appropriately designed institution can support high levels of cooperation. By contrast, under the classic model of public information information, which corresponds to an institution of size $Q = 1$, high rates of cooperation could only occur under Stern Judging. The public institution that maximizes cooperation for a given norm always requires more than $Q = 1$ members, although size $Q = 2$ is typically sufficient to achieve as much cooperation as possible under yet larger institutions.

As predicted by our analysis of replicator dynamics (Figure 2), we find that in finite populations governed by Stern Judging rates of cooperation are highest for strict institutions, which engender as much cooperation as empathetic societies using private moral assessment (Figure 3). Under the Simple Standing norm, we also find that strict institutional assessment performs best, and it sustains cooperation at levels that even exceed those under empathetic private assessment. By contrast, forgiving institutions produce less cooperation under Simple Standing or Stern Judging, although those institutions still outperform unempathetic private assessment (Figure 3).

Scoring is the only norm of moral evaluation that does not account for the recipient's reputation: it universally treats cooperation as a moral good. We find that the cooperation in societies governed by Scoring are nonetheless highly sensitive to the type of institution. Cooperation rates are highest under tolerant institutional assessment, where they greatly exceed the rates of cooperation in populations with private moral assessment (Figure 3). We observe a similar trend under Shunning, the norm that condemns any social interaction with individuals of bad moral character. In striking contrast to prior studies of public information (which corresponds to an institution of size $Q = 1$), we find that Shunning can engender cooperation levels as high as those supported

by Stern Judging, provided the institution has at least two members and is sufficiently tolerant (Figure 3).

Why do some norms require tolerant institutionalized moral assessment, while others require strict institutions to promote cooperation? The answer has to do with the norm’s sensitivity to occasional anti-social behavior. Shunning and Scoring are the two norms for which a single interaction with a bad individual can trigger a cascade of punishment and defection that eventually leads to low cooperation rates. In these cases tolerant institutions help individuals avoid being assigned bad reputations from occasional interactions with bad players, and so a high frequency of cooperating discriminators is less likely to be dislodged by occasional errors, mutation, or drift under tolerant Shunning or Scoring. Under Simple Standing, on the other hand, strict institutions promote more cooperation because only those individuals who defect against good are labeled bad, and high values of q help keep these defectors in check. The same mechanism holds under Stern Judging, where strict institutions are more efficient at inducing selection against willful defectors.

Evolution of institutional adherence

We have shown that high levels of cooperation can be achieved under second-order social norms in a society that uniformly follows an institution of public monitoring, with appropriate size and strictness. But the ancestral state of most human societies probably lacked any such institution or any instinct to adhere to a public broadcast. And so, initially at least, each individual was independently responsible for evaluating social interactions observed from their own perspective. A critical question remains: why would individuals choose to abandon their personal beliefs in favor of blind trust in the reputations reported by a centralized monitoring system?

To investigate whether trust in public monitoring can evolve via individual-level selection, we studied institutional adherence as a cultural trait that can itself spread from individual to individual by copying behavior, where individuals tend to copy the traits of others who have relatively higher payoffs. We developed replicator dynamics to quantify the selection pressure on institutional adherents competing against individuals who use private monitoring (Methods Section 2). We compared these predictions to Monte Carlo simulations in a finite population of DISC players, where we tested whether a single institution-following mutant can invade and overtake the population by social contagion, resulting in a population that uniformly adheres to institutional broadcasts.

We find that adherence to public institutions can often be favored by cultural evolution and eventually fix in the population. The spread of institutional adherence tends to homogenize the reputation matrix and eliminate unjustified defections that occur under private assessment of reputations [Radzvilavicius et al., 2019], thereby increasing individual-level fitness. The conditions that promote public monitoring depend upon the social norm, and they also depend on the level of empathy employed by those individuals who rely on private judgements (Figure 4).

In Stern Judging societies, adherence to a strict institution is much more likely to evolve than adherence to a tolerant institution; and adherence is favored to fix only if the ancestral population has low levels of empathetic perspective taking (Figure 4). Likewise, for Simple Standing, highly empathetic societies are not likely to evolve adherence to public monitoring, whereas populations dominated by non-empathetic moral evaluators are likely to evolve adherence to institutions with large Q , for all but the strictest values of the threshold q .

Under Scoring, institutions of size $Q = 1$ are neutral and hence unlikely to fix when rare. But adherence to a tolerant institution consisting of more than one observer is overwhelmingly favored by selection (Figure 4), and will eventually engender high levels of cooperation (Figure 3). This result is independent of empathy, since assessment by Scoring takes into account only the

Figure 4: The evolution of adherence to a public institution of moral assessment. We introduce a single institutional adherent (DISC-ADHERE) into a population consisting of non-adherents (DISC-PRIVATE), and we record the proportion of Monte Carlo simulations in which the institutional adherent eventually fixes in the population. Fixation probabilities exceeding $1/N = 1/50$ indicate that adherence to the institution is favored by selection, compared to neutral drift. Vertical dotted lines represent mean frequencies of “good” reputations in the population at the time of mutant introduction, averaged over all subjective perspectives. Horizontal dotted lines indicate fixation probabilities of neutral mutants. Error bars indicate the 95% confidence intervals for the SEM. Population size is $N = 50$, and error rates $e_1 = e_2 = 0.02$. The institution is comprised of an external board of Q members (see Methods Section 2, and also Supplementary Figure 7 for analogous results with an internal board).

action of the donor and not the reputation of her partner. On the other hand, societies with low empathy are unlikely to evolve institutional adherence under Shunning, as fixation probabilities of the mutant are always close to neutral. In empathetic societies of Shunners, however, tolerant institutions are favored by selection, resulting in fixation and high rates of cooperation.

Why are some types of institutions more likely to evolve widespread adherence than others? The equilibrium frequency of good individuals prior to the emergence of an institution (vertical dotted lines in Figure 4) is responsible for these differences. When individuals move from private to public monitoring, they move from using relative reputations to objective reputations (a positive effect): however, depending on the tolerance threshold q , they may lose some capacity to discriminate, which can lead to either unconditional cooperation or defection. Depending on the norm, then, those who follow an institution can sometimes be punished (assigned a bad reputation) by the majority of the society that does not yet follow the institution.

In cases where the equilibrium frequency of ‘good’ players is initially low, such as societies governed by Scoring and Shunning, high values of the institutional threshold q means that adherents will universally see all of their peers as bad individuals. As a result, an institution-following mutation will not invade, since under both of these norms an individual who defects against a bad recipient is punished, in effect, by being assigned a bad reputation. Those who do not choose

to trust the institutionalized reputation system will be better able to discriminate between the members of the society who are worthy of their help and the unworthy ones. And so only tolerant institutions, with q close to the mean frequency of “good” players in the ancestral population, will be able to invade under the norms of Scoring or Shunning (Figure 4).

On the other hand, under Stern Judging, where the initial frequency of good players is intermediate (e.g. 50% with $E = 0$), a rare mutant who adheres to a tolerant institution will see most players as good. As a result, the invader will engage in a high frequency of un-reciprocated cooperation and, in addition, she will be punished half of the time (assigned a bad reputation) by those peers who do not yet follow the institution. Because of this double cost, rare mutants who adhere to the institution cannot invade. With high values of q , on the other hand, when the institution-following invader is rare most of the reputations provided by the institution will be bad, especially when Q is large. As a result, the invader will not cooperate, she will receive a higher-than-average payoff, and she will be assigned a bad reputation only half of the time. As the invader frequency increases, so does the frequency of good players (averaged over all perspectives), and institutional adherence will eventually fix. Similar dynamics apply to a rare institution-following invader under the Simple Standing norm, but in this case the frequency of good players is already high in the ancestral population and cooperation is rewarded from the very beginning of the invasion.

Whatever the social norm, once the trait responsible for following a public institution fixes in a population, it is thereafter robust. That is, adherence cannot be selectively displaced by a mutant that uses only private moral assessment (data not shown). This means that institutions of moral assessment are evolutionarily stable. There is a simple intuition that explains this stability. Once the vast majority of a population follows the public institution, the rare invader who relies on her own subjective assessment is at a disadvantage because she may hold different beliefs about reputations than others hold. For instance, the rare invader will sometimes see an institutionally good individual as bad, will defect, and will subsequently be punished through reciprocal defection by the institution-following majority.

To illustrate the frequency-dependent effects summarized above, we calculated the mean fixation probability of adherence, starting from a random initial frequency of institution-following mutants drawn uniformly from the open interval $(0, 1)$. The mean fixation probability over this distribution of initial conditions provides a measure of the stability of institutional adherence against those who use private reputations. The mean fixation probability depends, once again, on the type of the institution (q and Q), the social norm, and the degree of empathy among non-adherents (Supplementary Figure 1). Under Stern Judging and low empathy, strict institutions for public monitoring have a high fixation probability over almost the entire range of initial mutant frequencies, resulting in a high mean fixation probability. Under all other norms, more tolerant institutions are stabilized by evolution. Similar to our observations about the fate of new mutant (Figure 4), we find that higher rates of empathy tend to reduce the mean fixation chance of adherence under Stern Judging and Simple Standing (Supplementary Figure 1).

The simulations above describe institutional adherents competing against private assessors, assuming that everyone pays attention to reputations, i.e., assuming DISC players. We have also explored the fate of discriminators who adhere to public broadcasts (henceforth DISC-ADHERE) when competing against more diverse alternatives: discriminators using private assessment (DISC-PRIVATE), unconditional defectors (ALLD), and unconditional cooperators (ALLC). Even when introduced into a population containing these three other types, institutional adherents are often favored by selection, although their fixation probability may be reduced (Supplementary Figures 3-6), especially when ALLD is frequent. Still, for some norms and institutions, such as strict Stern Judging, selection can promote adherence even in a population initially dominated by ALLD (Supplementary Figure 5) or by ALLC (Supplementary Figure 6).

Crowding effects of empathy and institutions

Whether evolution promotes adherence to a public institution of judgement depends on whether private assessors have an internal capacity for empathy. We have seen that adherence is less likely to spread in a population of empathetic private assessors, as opposed to egocentric assessors (Figure 4). This result makes intuitive sense, because empathy allows populations to achieve high levels of cooperation even in the absence of a public information [Radzvilavicius et al., 2019], and so there is less marginal benefit for empathetic individuals to adhere to a public institution.

These results mirror the effect known as motivational crowding, which has broad theoretical and empirical support in economics [Frey, 1994, Frey and Oberholzer-Gee, 1997, Frey and Jegen, 2001]: external interventions can under undermine intrinsic motivations. It is interesting to note, however, that the general pattern of interference between intrinsic empathy and extrinsic institutions is not ubiquitous. Figure 4 reveals two counter-examples: first, adherence to a strict public institution can selectively invade a population under the Stern Judging norm, whether or not the initial population is empathetic; and second, adherence to a large tolerant institution can invade under the Shunning norm, even when the private assessors are empathetic. (Note that in the case of the Scoring norm, empathy cannot possibly interfere with institutional adherence, because this social norm does not account for the recipient’s reputation when assessing a donor.) The extent to which these two counter-examples coincide with the few identifiable cases of synergy between external and internal incentives in the economics literature [Frey and Jegen, 2001] remains a topic for further study.

Replicator dynamics of institutional adherence

We have used Monte Carlo simulations to study whether institutional adherence will spread in a finite population (Figure 4). But these simulations do not provide a systematic account of how adherence depends on parameters, including error rates (e_1 and e_2), the costs and benefits cooperation (b and c), institution size (Q), and institution strictness (q). To address these questions analytically we developed replicator dynamic equations describing competition between institutional adherents and private assessors in an infinite population (Methods Section 3). These replicator dynamics are more complicated those that describe strategy competition among adherents alone, because in this setting two individuals may disagree about a third individual’s reputation.

We have derived replicator dynamics to describe competition among four types of individuals: unconditional cooperators (ALLC), unconditional defectors (ALLD), discriminators who adhere to the public institution (DISC-ADHERE), and discriminators who rely on their own private assessments (DISC-PRIVATE). This derivation requires that we separately track the frequency of individuals within each of these four types that is seen as good either by private assessors or by the public institution, as well as the frequency of disagreements both among private assessors and between private assessors and public adherents. Finally, we must specify whether the Q members of the institution are themselves engaged in strategic interactions (“internal board”, Methods Section 2), or whether the institutional assessors are separate from the population (“external board”).

The resulting replicator dynamics describe how all model parameters – including the error rates, payoff matrix, and design of the public institution – influence dynamics in an infinite population, for each of the four social norms we have studied. One insight this immediately provides is that selection favors adherence only if the benefit-to-cost ratio b/c exceeds a critical value that depends on the error rates, social norm, and institutional design (Methods Eq. 15).

Figure 5 shows the growth rate of a rare adherent type (DISC-ADHERE) in a population otherwise comprised of private assessors (DISC-PRIVATE). The benefit-to-cost ratio and the error

Figure 5: Growth rate of institutional adherents (DISC-ADHERE) in a population otherwise composed of private assessors (DISC-PRIVATE). The instantaneous growth rate per unit time is computed according to replicator dynamics in an infinite population (Methods Eq. 11) with the initial frequency of adherents set at $1/50$, for strict ($q > 1/2$) or tolerant ($q < 1/2$) institutions of size $Q = 2$ (external board). Varying the benefit-to-cost ratio can have a qualitative impact on the evolution of institutional adherence. For example, under tolerant Stern Judging, adherents are strongly suppressed by selection for $b/c < 50$ but favored by selection for $b/c > 50$. Unless otherwise varied in a figure panel, parameter values are $e_1 = e_2 = 0.02$, $b = 5$, $c = 1$.

rates can have substantial effects on the magnitude of selection for or against institutional adherents. Generally speaking, for both strict and tolerant institutions, variation in b/c tends to have larger effects on adherent growth compared to variation in errors of assessment (e_2) or errors of strategy implementation (e_1). These analyses uncover several striking qualitative phenomena: conditions where the benefit-to-cost ratio determines whether adherence is selectively favored or disfavored. Under tolerant Stern Judging, we have already seen that selection *disfavors* institutional adherents for $b/c = 5$ (Figure 4), in agreement with replicator-dynamic predictions (Figure 5); and yet replicator dynamics imply adherence to tolerant Stern Judging will be selectively *favored* for sufficiently large benefit-to-cost ratios (e.g. $b/c > 50$, Figure 5). We verified this prediction by Monte Carlo simulations in finite populations (Supplementary Figure 8). Likewise, under strict Shunning, replicator dynamics imply that selection will favor institutional adherence when $b/c < 50$, but selection will disfavor adherence for $b/c > 50$ (Figure 5); and under tolerant Shunning, replicator dynamics imply selection will favor adherence more strongly when b/c is larger (Figure 5). All of these predictions were verified by Monte Carlo simulations in finite populations (Figure 4 and Supplementary Figure 8). Similar results hold for an “internal board” of institution members (Supplementary

Figures 7, 9, and 13).

Discussion

Much of human prosocial behavior is guided by social norms and informal moral codes of fairness [Ostrom, 2000, Fehr et al., 2002, Fehr and Fischbacher, 2004]. Societies also adopt formal institutions that reinforce these informal constructs to enforce conditional cooperation and resolve conflict. For instance, local institutions help resolve public goods dilemmas in irrigation systems [Meinzen-Dick, 2007] and management of forest commons [Rustagi et al., 2010]. Likewise, e-commerce companies aggregate individual assessments of the reputations of buyers and sellers, providing a public broadcast to a large community of users Utz et al. [2009]. Credit bureaus synthesize and publicize the reputations of borrowers, so that lending agencies can choose to reward cooperative behavior with easy access to future capital McIntosh et al. [2013]. Analogous institutions even existed in the era of pre-state trade [Milgrom et al., 1990, Greif, 2004].

Public knowledge of reputations is the default assumption for the theory of cooperation based on indirect reciprocity, in which cooperation is conditioned on the reputations of potential exchange partners. In reality, however, reputations are often assigned by private monitoring of social interactions by independent individuals, and so they may be observer-dependent. Private monitoring and moral relativity – as recent theoretical studies show – causes the breakdown of cooperation [Hilbe et al., 2018, Okada et al., 2017, Radzvilavicius et al., 2019]. Because private monitoring was likely the ancestral state in early societies, institutionalized moral systems had to emerge by mutual agreement and adherence spread by social contagion.

We have analyzed how to design public institutions of monitoring that stimulate the highest levels of cooperation in a society, under the theory of indirect reciprocity. But even when such an option is available, the question remains: why would people choose to blindly trust moral judgments provided by an institutional observer, rather than trusting their own moral compass? To address this question we have also studied whether adherence to a public institution will spread by social contagion, in a population that initially relies on private judgements. Our key finding is that adherence to centralized institutions of assessment can often spread when rare, as individuals that abandon their private beliefs are on average more successful than their peers. If the entire population trusts a public monitoring system then the rate of disagreement about the moral standing of individuals is reduced, which eliminates unjustified punishment. Nonetheless, the specific form of public institution that provides the greatest benefit depends on the social norm. Larger institutions tend to outperform small ones, and they can perform far better than broadcasting the judgements of a single individual. Tolerant institutions stimulate more cooperation for certain social norms of moral judgement, whereas strict institutions do better for other norms.

Our results provide a theoretical justification for a key assumption in the theory of indirect reciprocity: institutionalized public monitoring. Critically, we also find that institutions can be designed to dramatically increase cooperation rates, even for social norms that previous studies had found to perform poorly, such as Scoring, Shunning, and Simple Standing. Cooperation requires choosing an institution of the appropriate size and tolerance, given the prevailing social norm. For example, societies that adhere to Scoring or Shunning norms can achieve rates of cooperation as high as those found in societies governed by Stern Judging, provided the centralized monitoring system consists of more than a single individual and is highly tolerant. These results may have practical implications for improving the design of rating and reputation systems in e-commerce and sharing economies.

A different solution to the problem of private monitoring, and consequent breakdown in

cooperation, is empathetic perspective taking [Radzvilavicius et al., 2019]. But it is difficult to know how effective empathy is at promoting cooperation in the face of cognitive costs and the potential for cheating or manipulation. Nonetheless, once individuals have acquired a capacity for empathy, this internal mechanism often “crowds out” selection that would otherwise favor adherence to an public institution of centralized monitoring. The extent of interference between these internal and external mechanisms of moral assessment, as well as the precise relationship to motivation crowding in the economics literature, remains a topic for future research.

Our study has not addressed how an institution of public monitoring is established to begin with. Instead, we have focused on whether, and when, adherence to an existing public broadcast is beneficial. The question of establishing an institution *de novo* is perhaps more difficult – and more properly a question in political science than in evolutionary game theory. In practice, public institutions of judgement are often formed through governance, by mutual agreement in a constitution, or by legislation. Although the coordination problem of forming an institution is outside the scope of our study, we can nonetheless provide some rudimentary insights by quantifying the level of taxation that rational individuals would willingly pay to support such an institution. In the context of the donation game we have studied, a well-designed institution can be expected to replace a defecting population of private assessors with a cooperative population of institutional adherents, increasing the per capita payoff by roughly $b - c$. And so, in principle, each individual in the population would be willing (or indifferent) to contribute up to $b - c$ per round to support the formation and operation of such a public institution.

Our analysis has also neglected the problem of institutional corruption. For example, what is to stop a member of the institution from engaging in side deals with individuals who seek to extract a good reputation broadcast in exchange for a bribe? Here again we can at least make a rudimentary analysis of the incentives and costs associated with corruption. Each member of the institution can expect to receive an income of roughly $N(b - c)/Q$ per round, as described above, in a society of N rational individuals who understand the value of public monitoring. And so, even in the simple one-shot setting, if corruption carries a chance of discovery p that is penalized by loss of income, then an institution member will be robust against a bribe B provided $(1 - p)B < N(b - c)/Q$. This simple calculus provides approximate bounds on the degree of wealth inequity that can be tolerated while maintaining an incorruptible institution of public assessment.

Although the theory of indirect reciprocity is limited to pairwise interactions, the simple idea of behavior conditioned on partner reputation is a compelling explanation for cooperation. The role of reputations in n -player interactions, such as in public goods games, remains largely unexplored. Likewise, the concept of a social norm in this context – that is, the moral heuristic that determines which types of actions result in assigning someone a good or bad reputation – require further theoretical and empirical study. Very little is known about the origin and spread of such norms. The standard assumption in the literature is that norms are universally accepted or imposed externally and do not change over time. There a few studies on norm competition [Uchida and Sigmund, 2010, Yamamoto et al., 2017], as well as ongoing discovery of new, effective social norms [Sasaki et al., 2017]. These studies have begun to address how the definition of moral good may itself evolve, via random exploration and social contagion. But it remains unclear which social norms of moral judgement will emerge as norms are allowed to evolve from personal beliefs in societies, either with or without public institutions of assessment. We also lack a mechanistic theory for how the definitions of moral codes are transmitted among individuals, or whether any *a priori* restrictions limit the norms that individuals can choose or transmit. These questions await future research.

Methods

1 Strategy evolution under public monitoring

We start by analyzing the evolutionary dynamics of behavioral strategies in an infinite population, assuming complete adherence to a public institution of moral assessment. We use classical replicator dynamics [Taylor and Jonker, 1978] to describe how the frequencies of three strategic types change in time. We consider the following strategies: always cooperate (ALLC, denoted X), always defect (ALLD, denoted Y), and discriminate (DISC, denoted Z). Unlike ALLC or ALLD, the DISC type conditions behavior on reputations: DISC cooperates with individuals that have a good reputation and defects with individuals that have a bad reputation.

1.1 Replicator dynamics

The frequencies of the three strategic types, f_X , f_Y , and f_Z change according to their fitness relative to the mean population fitness:

$$\begin{aligned} \dot{f}_i &= f_i(\Pi_i - \bar{\Pi}), \text{ with} \\ \bar{\Pi} &= \sum_{j \in \{X, Y, Z\}} f_j \Pi_j \end{aligned} \quad (1)$$

where $i \in \{X, Y, Z\}$ denotes the i th strategic type and Π_i denotes the mean fitness of type i .

The fitness of each individual is determined by interactions in the donation game with all other individuals in the population. Each individual interacts once as donor and once as recipient with all other individuals in the population. Each individual derives a benefit b from interaction with every cooperator and from interaction with a every discriminator who views them as good (i.e. if the focal individual has a good institutional reputation). Cooperators always pay the cost c of cooperation, and discriminators pay the cost only to individuals who have a good reputation; defectors never pay the cost. Finally, with probability e_1 , an individual who intends to cooperate will accidentally defect. As a result, the mean fitness of individuals with each strategic type is given by

$$\begin{aligned} \Pi_X &= (1 - e_1)b(f_X + f_Z G_X) - (1 - e_1)c \\ \Pi_Y &= (1 - e_1)b(f_X + f_Z G_Y) \\ \Pi_Z &= (1 - e_1)b(f_X + f_Z G_Z) - (1 - e_1)Gc, \end{aligned} \quad (2)$$

where G_i denotes the fraction of individuals of type i who have good institutional reputations and $G = \sum_i f_i G_i$ denotes the total fraction of the population with a good institutional reputation.

1.2 Reputation dynamics

We assume that reputation frequencies reach equilibrium before strategy frequencies change – that is, we assume the timescale of reputation updating is fast compared to that of strategic updates Okada et al. [2018]. In this section we describe the equilibrium reputations that emerge, assuming complete adherence to the public institution of moral assessment.

Reputations in our model are publicly known and determined by the consensus view by provided by an “institution”. The institution consists of Q members. Each round, each institution member observes a independent, random interaction of each individual acting as a donor. As in

Okada et al. [2018], we assume that each observer makes her moral evaluation of a donor’s action based on an independently observed social interaction.

An institution member updates the reputation of a donor according to her action (cooperate or defect) towards a recipient. The resulting reputation that member ascribes to the donor is governed by the social norm of the population. After forming their individual assessments, if at least a fraction q of the institution members agree that an individual’s reputation is good, then the individual’s reputation is broadcast as good to the entire population. Otherwise, it is broadcast as bad. Table 1 summarizes the probability that an institutional observer will assess a donor’s reputation as good, depending on the observer’s view of the recipient and the donor’s intended action and the social norm. These probabilities account for the two types of errors that occur in our model: a donor’s cooperative intention is erroneously executed as defection with probability e_1 , and an observer erroneously assigns a bad reputation instead of a good reputation, and *vice versa*, with probability e_2 .

observer view of recipient	donor intent	P_{SJ}	P_{SS}	P_{SC}	P_{SH}
good	cooperate	ϵ	ϵ	ϵ	ϵ
good	defect	e_2	e_2	e_2	e_2
bad	cooperate	$1 - \epsilon$	$1 - e_2$	ϵ	e_2
bad	defect	$1 - e_2$	$1 - e_2$	e_2	e_2

Table 1: Probability P_k that an observer will assign a donor a good reputation, under each of four social norms $k \in [\text{stern judging, simple standing, scoring, shunning}]$. Here, $\epsilon = (1 - e_2)(1 - e_1) + e_2e_1$ is the probability that an individual who intends to cooperate with a recipient who has a good reputation is ultimately themselves assigned a good reputation – either because they successfully cooperate and are correctly assigned a good reputation (first term) or accidentally defect and are erroneously assigned a good reputation (second term).

The probabilities in Table 1 will help us to derive the equilibrium proportion of individuals of type i that are viewed as having a good reputation in the eyes of an arbitrary observer, which we denote g_i . The views of the Q observers that comprise the institution will then collectively determine the institutional broadcasts, and thus G_i , as described below.

The equilibrium proportion of type i that are assigned a good reputation by an arbitrary observer can be expressed as the solution to a system of three simultaneous equations, one for each $i \in \{X, Y, Z\}$. We derive these equations in detail for the Stern Judging norm, referring to Table 1 as needed. (Analogous derivations apply for the three other norms.) Consider first the chance, g_X , that an unconditional cooperator will be assigned a good reputation by a observer. An unconditional cooperator ends up with a good reputation in the following two distinct ways:

1. She interacts with someone who has a good reputation (probability G), intending to cooperate, and is successfully assigned a good reputation by the observer (probability ϵ , from Table 1).
2. She interacts with someone who has a bad reputation (probability $1 - G$), intending to cooperate, and is erroneously being assigned a good reputation (probability $1 - \epsilon$, from Table 1).

Likewise, an unconditional defector ends up with a good reputation in two ways:

1. She interacts with someone who has a good reputation (probability G), intending to defect, and is erroneously assigned a good reputation (probability e_2 , from Table 1).

2. She interacts with someone who has a bad reputation (probability $1 - G$), intending to defect, and she is correctly assigned a good reputation (probability $1 - e_2$, from Table 1).

Finally, a discriminator ends up with a good reputation in two ways:

1. She interacts with someone who has a good reputation (probability G), intending to cooperate, and is correctly assigned a good reputation (probability ϵ , from Table 1).
2. She interacts with someone who has a bad reputation (probability $1 - G$), intending to defect, and is correctly assigned a good reputation (probability $1 - e_2$, from Table 1).

In summary, for a population following the Stern Judging norm, the equilibrium frequency of good individuals in each strategic type from the perspective of an arbitrary observer satisfy the following equations:

$$\begin{aligned} g_X &= \epsilon G + (1 - \epsilon)(1 - G) \\ g_Y &= e_2 G + (1 - e_2)(1 - G) \\ g_Z &= \epsilon G + (1 - e_2)(1 - G). \end{aligned} \tag{3}$$

An analogous derivation for Simple Standing produces

$$\begin{aligned} g_X &= \epsilon G + (1 - e_2)(1 - G) \\ g_Y &= e_2 G + (1 - e_2)(1 - G) \\ g_Z &= \epsilon G + (1 - e_2)(1 - G). \end{aligned} \tag{4}$$

Likewise, for Scoring we have

$$\begin{aligned} g_X &= \epsilon \\ g_Y &= e_2 \\ g_Z &= \epsilon G + (1 - e_2)(1 - G). \end{aligned} \tag{5}$$

Finally, for Shunning we have

$$\begin{aligned} g_X &= \epsilon G + e_2(1 - G) \\ g_Y &= e_2 \\ g_Z &= \epsilon G + e_2(1 - G). \end{aligned} \tag{6}$$

These equations correspond to the expressions under private assessment with complete empathy from Radzvilavicius et al. [2019], with the exception that $g = \sum_i f_i g_i$, the proportion of the population with a good reputation in the eyes of an arbitrary observer, is replaced here by $G = \sum_i f_i G_i$, the proportion of the population with a good institutional reputation.

For each social norm, the system of equations above is closed once we specify how G depends on g_X , g_Y , g_Z and on the design of the institution. We consider two limiting cases: very strict institutions that broadcast an individual as good only if all members agree she is good ($q > (Q - 1)/Q$), and very tolerant institutions that broadcast an individual as good provided at least one member views her as good ($q < 1/Q$). In these two respective cases we have

$$G_i = \begin{cases} 1 - (1 - g_i)^Q & \text{for tolerant institutions, } i \in \{X, Y, Z\}, \\ g_i^Q & \text{for strict institutions, } i \in \{X, Y, Z\}. \end{cases} \tag{7}$$

For completeness' sake, we provide the expression for arbitrary q and Q :

$$G_i = \sum_{k=\text{ceil}(qQ)}^Q \binom{Q}{k} g_i^k (1 - g_i)^{Q-k}. \quad (8)$$

The resulting system of equations for g_X, g_Y , and g_Z can be solved by radicals when $Q \leq 2$. More generally, a unique feasible solution exists for any Q , and it can be computed numerically by iterating the above system of equations, after choosing any initial value $(g_X^0, g_Y^0, g_Z^0) \in (0, 1)^3$. Successive iterates (g_X^k, g_Y^k, g_Z^k) remain in $[0, 1]^3$ and form a contraction: the equations for g_X, g_Y, g_Z are each convex combinations of elements in $(0, 1)$, provided $e_2 > 0$ and $e_1 > 0$. Solving this system determines the equilibrium frequency of good individuals of each strategic type from the eyes of an arbitrary observer (g_i) and therefore also yields the institutional broadcast (G_i , from Eq. 7, and $G = \sum_i f_i G_i$). Substituting these expressions into the replicator equation (Eq. 1 and Eq. 2) provides the dynamics of strategy frequencies f_i , allowing us to compute selection gradients, strategic equilibria, and basins of attractions, as shown in Figure 2.

2 Competition between institutional adherents and non-adherents

Whereas the analysis in Section 1 assumes complete adherence to the reputations broadcast by the public institution, in this section we now accommodate a mixture of individuals, some of whom condition their behavior based on their own private assessment of recipient reputations, and others who condition their behavior based on the institutional broadcast. This distinction has no effect on the behavior of unconditional cooperators or defectors (ALLC or ALLD), but it does affect the behavior of discriminators. And so, henceforth we divide the discriminators into two types: *adherents* (DISC-ADHERE, denoted Z_a) and *non-adherents* (DISC-PRIVATE, denoted Z_n).

Adherents use the reputations broadcast by the institution when determining whether to cooperate or defect with an individual. Non-adherents do not: they use their private view of the individual, which was formed by observing a random interaction in which that individual acted as a donor. Here we develop replicator dynamics to describe competition between these two types of discriminators, adherents Z_a and non-adherents Z_n . (Eventually we extend this to describe competition between all four types, X, Y, Z_a , and Z_n .)

We consider two alternative assessment protocols for the *institution members* themselves. (This distinction was not relevant in the context of complete adherence, in the preceding section.) In particular, we consider two possibilities for the institution:

1. The institution may be comprised of an *internal board* of Q individuals sampled from the population. In this case the internal board may consist of adherents, non-adherents, or a mix thereof. When a board member is an adherent, they will use the institutional assessment of a recipient when updating their opinion of a donor's reputation. When a board member is a non-adherent, they will use their own private view of a recipient, which may differ from the public institutional one, when assessing a donor.
2. The institution may alternatively be comprised of an *external board* of Q individuals who are all adherents, irrespective of the strategy composition of the population. External board members rely solely on the institutional public broadcast of the recipient when formulating their opinion of a donor's reputation. We describe this configuration as "external" because the institution members act "outside" of the population, using institutional assessments to determine reputations irrespective of the actual composition of the population. Equivalently,

the external board can be interpreted as an institution whose members are inside the population itself, provided the institution members diligently compartmentalize their own personal reputation assessments (when acting as a donor) from the institutional assessment (when updating their views of a recipient for contribution to the institution).

2.1 Replicator dynamics

Because adherents and non-adherents may not cooperate and defect in the same situations, their reputation dynamics must be considered separately. Let $g_{i,j}$ denote the proportion of individuals following strategy i that are seen as having a good reputation by an individual with strategy j . As before we let G_i denote the proportion of the individuals of type i that have a good institutional reputation. Define

$$\begin{aligned} G &= \sum_i f_i G_i, \\ g_\bullet &= \sum_i f_i g_{i,Z_n} \\ \gamma_2 &= \sum_i f_i G_i g_{i,Z_n}, \\ g_2 &= \sum_i f_i g_{i,Z_n}^2. \end{aligned} \tag{9}$$

Here, G is (as previously) the proportion of the total population with a good institutional reputation, $g_\bullet$ is the proportion of the total population that a non-adherent sees as having a good reputation, γ_2 is the probability that an institutional adherent and non-adherent agree that a given individual's reputation is good, and g_2 is the probability that two non-adherents agree that a given individual's reputation is good. The strategic types i in Eq. 9 range over Z_a and Z_n in this section (but they will additionally include X and Y in Section 2.3).

The mean fitness of individuals of type Z_n and Z_f can now be written as

$$\begin{aligned} \Pi_{Z_n} &= (1 - e_1)b(f_{Z_n}g_{Z_n,Z_n} + f_{Z_a}G_{Z_n}) - c(1 - e_1)g_\bullet \\ \Pi_{Z_a} &= (1 - e_1)b(f_{Z_n}g_{Z_a,Z_n} + f_{Z_a}G_{Z_a}) - c(1 - e_1)G, \end{aligned} \tag{10}$$

which notably differ from the expressions in the context of complete adherence. Finally, the replicator dynamics are given by

$$\begin{aligned} \dot{f}_i &= f_i(\Pi_i - \bar{\Pi}), \text{ with} \\ \bar{\Pi} &= \sum_{j \in \{Z_n, Z_a\}} f_j \Pi_j. \end{aligned} \tag{11}$$

Note that g_{Z_n, Z_a} and g_{Z_a, Z_n} do not appear in the fitness expressions or replicator dynamic equations. This is because adherents' private assessments of reputations do not directly determine their behavior: their behavior is governed solely by institutional assessments.

2.2 Reputation dynamics

Reputation dynamics and a population containing both adherents (Z_a) and non-adherents (Z_n) must account for all possible configurations of an observer assessing a donor interacting with a recipient (Table 2). By cross-referencing Table 2 with Table 1 provides expressions for $g_{i,j}$, the proportion of individuals following strategy i that are seen as good by an individual of strategy j .

observer strategy	observer view of recipient	donor strategy	donor view of recipient	probability of configuration	donor intent
Z_a	good	Z_a	good	G	cooperate
Z_a	good	Z_a	bad	0	defect
Z_a	bad	Z_a	good	0	cooperate
Z_a	bad	Z_a	bad	$1 - G$	defect
Z_a	good	Z_n	good	γ_2	cooperate
Z_a	good	Z_n	bad	$G - \gamma_2$	defect
Z_a	bad	Z_n	good	$g_\bullet - \gamma_2$	cooperate
Z_a	bad	Z_n	bad	$1 - G - g_\bullet + \gamma_2$	defect
Z_n	good	Z_a	good	γ_2	cooperate
Z_n	good	Z_a	bad	$g_\bullet - \gamma_2$	defect
Z_n	bad	Z_a	good	$G - \gamma_2$	cooperate
Z_n	bad	Z_a	bad	$1 - G - g_\bullet + \gamma_2$	defect
Z_n	good	Z_n	good	g_2	cooperate
Z_n	good	Z_n	bad	$g_\bullet - g_2$	defect
Z_n	bad	Z_n	good	$g_\bullet - g_2$	cooperate
Z_n	bad	Z_n	bad	$1 - 2g_\bullet + \gamma_2$	defect

Table 2: Possible configurations of an observer, donor, and recipient. The fifth column indicates the probability of each configuration, and the sixth column indicates the donor's intent (which follows directly from the donor's view of the recipient). Cross-referencing the fifth and sixth columns with Table 1 yields the probability that the donor will be assessed as good in the eyes of the observer.

For example, under the Stern Judging norm, we have the following system of equations for equilibrium values of $g_{i,j}$:

$$\begin{aligned}
g_{Z_a, Z_a} &= G\epsilon + (1 - G)(1 - e_2), \\
g_{Z_a, Z_n} &= \gamma_2\epsilon + (g_\bullet - \gamma_2)e_2 + (G - \gamma_2)(1 - \epsilon) + (1 - g_\bullet - G + \gamma_2)(1 - e_2), \\
g_{Z_n, Z_a} &= \gamma_2\epsilon + (g_\bullet - \gamma_2)(1 - \epsilon) + (G - \gamma_2)e_2 + (1 - g_\bullet - G + \gamma_2)(1 - e_2), \\
g_{Z_n, Z_n} &= g_2\epsilon + (g_\bullet - g_2)(e_2 + 1 - \epsilon) + (1 - 2g_\bullet + g_2)(1 - e_2).
\end{aligned} \tag{12}$$

Note that for a population consisting entirely of adherents, these equations reduce to the case of non-empathetic private assessment that has been studied previously Radzvilavicius et al. [2019]. Analogous equations for Simple Standing, Shunning, and Scoring can be immediately obtained in the same manner.

The system of equations for equilibrium $g_{i,j}$ is closed once we specify how individual views are aggregated to produce institutional assessments, G_i . In the case of an external board whose members always make assessments based on the institutional view of a recipient, we have the following:

$$G_i = \begin{cases} 1 - (1 - g_{i, Z_a})^Q & \text{for tolerant institutions, external board,} \\ g_{i, Z_a}^Q & \text{for strict institutions, external board.} \end{cases} \tag{13}$$

For the case of an internal board, the composition of adherents and non-adherents among institution members is determined by randomly drawing from the general population. And so in this case we

have

$$G_i = \begin{cases} \sum_{k=0}^Q \binom{Q}{k} f_{Z_a}^{Q-k} f_{Z_n}^k [1 - (1 - g_{i,Z_a})^{Q-k} (1 - g_{i,Z_n})^k] & \text{for tolerant institutions, internal board,} \\ \sum_{k=0}^Q \binom{Q}{k} f_{Z_a}^{Q-k} f_{Z_n}^k g_{i,Z_a}^{Q-k} g_{i,Z_n}^k & \text{for strict institutions, internal board.} \end{cases} \quad (14)$$

In either case, we can substitute the expressions for G_i and $G = \sum_i f_i G_i$ into Eq. 12 to obtain a closed system of equations for the equilibrium values of g_{ij} . For any Q , this system of equations has a unique solution in $[0, 1]^4$, which can be constructed by numerical iteration, because the right-hand-side of each equation is a convex combination of elements in $(0, 1)$. The equilibrium reputations can then be substituted into the replicator equation (Eq. 11) to describe the dynamics of strategy frequencies f_i over time.

Note that the condition for selection to favor (positive) growth of adherents, Z_a , is given by

$$\begin{aligned} 0 &< \Pi_{Z_a} - \Pi_{Z_n} \\ \therefore 0 &< b(f_{Z_n}[g_{Z_a,Z_n} - g_{Z_n,Z_n}] + f_{Z_a}[G_{Z_a} - G_{Z_n}]) - c(G - g_\bullet) \\ \therefore \frac{b}{c} &> \frac{G - g_\bullet}{f_{Z_n}[g_{Z_a,Z_n} - g_{Z_n,Z_n}] + f_{Z_a}[G_{Z_a} - G_{Z_n}]} \end{aligned} \quad (15)$$

The only parameters that appear directly in this expression are the game payoffs b and c , whereas e_2 and e_1 are involved in the expressions for reputations on the right-hand side. And so we see that selection favors adherence only if the benefit-to-cost ratio b/c exceeds a critical value that depends on the error rates, social norm, and institutional design.

2.3 Multiple strategic types

Finally, we can extend this analysis to include cooperators, defectors, and both flavors of discriminators. As before, we differentiate between the reputation assessments of institution adherents and non-adherents. Under the Stern Judging norm, the reputations of cooperators and defectors will be given by

$$\begin{aligned} g_{X,Z_a} &= G\epsilon + (1 - G)(1 - \epsilon), \\ g_{X,Z_n} &= g_\bullet\epsilon + (1 - g_\bullet)(1 - \epsilon), \\ g_{Y,Z_a} &= Ge_2 + (1 - G)(1 - e_2), \\ g_{Y,Z_n} &= g_\bullet\epsilon + (1 - g_\bullet)(1 - e_2). \end{aligned} \quad (16)$$

Reputations under other norms can be derived analogously. The replicator dynamics can be extended from equation 10:

$$\begin{aligned} \Pi_X &= (1 - e_1)b(f_X + f_{Z_n}g_{X,Z_n} + f_{Z_a}G_X) - c(1 - e_1) \\ \Pi_Y &= (1 - e_1)b(f_X + f_{Z_n}g_{Y,Z_n} + f_{Z_a}G_Y) \\ \Pi_{Z_n} &= (1 - e_1)b(f_X + f_{Z_n}g_{Z_n,Z_n} + f_{Z_a}G_{Z_n}) - c(1 - e_1)g_\bullet \\ \Pi_{Z_a} &= (1 - e_1)b(f_X + f_{Z_n}g_{Z_a,Z_n} + f_{Z_a}G_{Z_a}) - c(1 - e_1)G. \end{aligned} \quad (17)$$

3 Stochastic simulations

To study cooperation in finite populations of institution adherents, we performed a series of Monte Carlo simulations implemented in Julia. Each population consists of $N = 50$ individuals, all following the same social norm but each with its own strategy: ALLC, ALLD, or DISC. In each

successive discrete generation, every individual interacts with every other individual in three roles: once as a donor, once as a recipient, and once as an observer.

First, each individual plays a single round of the donation game with every member of the population, taking actions according to their own strategy and the (public) reputation of the recipient. We include self-interactions, although individuals cannot derive any net benefit from self-interactions. Individuals who intend to cooperate will accidentally defect with probability e_1 . Each individual pays a cost c for every game in which they cooperate and accrues a benefit b from every game in which their interaction partner cooperates.

The institution of public monitoring consists of $1 \leq Q \leq N$ designated players, each with their own subjective view (either “good” or “bad”) of the reputations of the rest of the population. The institution members remain the same throughout a given simulation. For every member i of the population, each institution member observes a single random interaction in which i acted as a donor. They then update their opinion of i based on i ’s action, the public reputation of the recipient j , and the social norm. With probability e_2 they assign the wrong reputation. If more than a threshold number qQ of institution members consider an individual to be “good”, that individual’s reputation is publicly broadcast as “good”. Otherwise, it is publicly broadcast as “bad”. These updated public reputations are used by the entire population in the following round to determine their behavior.

Selection and drift are implemented as a social contagion process using pairwise comparison. After the reputation updating step, a random pair of individuals i and j is chosen. Player i adopts the strategy of j with probability $1/(1 + \exp(-w[\Pi_i - \Pi_j]))$, where w is a selection strength parameter and Π_i and Π_j are the average payoffs earned by i and j (that is, the total payoff divided by N interactions). We set $w = 1$ for all simulations. Finally, with probability $\mu = 0.025$, an individual is randomly chosen to switch (or “mutate”) to a random strategy, following Santos et al. [2016].

We initialize each simulation with random strategies and random reputations. We then record the mean cooperation rate sampled at the end of 10,000 generations, averaged over 2,500 independent replicate simulations. We compared our results to simulations in which individuals rely exclusively on their own private reputation assessments but may, with probability E , empathetically adopt the perspective of the donor in each interaction as described by Radzvilavicius et al. [2019]. The results of these simulations were used to produce Figure 3.

Next, to study the evolution of adherence to a public institution, we adjusted the simulation to account for a mix of adherents and non-adherents. We initialized each population to consist solely of discriminators but added a binary trait corresponding to whether an individual uses their own private reputation assessments (non-adherent, Z_n) or the institutionally broadcast public assessments (adherent, Z_a). (Mutation is set to zero in these simulations). Each population initially consists entirely of non-adherents: after an equilibration period (250 generations), a single adherent is introduced in the population. In each subsequent generation, a random pair of individuals i and j is selected, and i adopts j ’s type (either Z_a or Z_n) with probability $1/(1 + \exp(-w[\Pi_i - \Pi_j]))$, as above. In addition, we either mandate that all Q institution members be adherents (external board) or allow them to use the assessment strategy of the first Q individuals in the population (internal board). Finally, we permitted empathetic evaluation with probability E as outlined in Radzvilavicius et al. [2019]. These simulations continued until the institution following trait either fixes or goes extinct. We record the proportion of fixations observed among 2,500 replicate simulations per parameter combination. These simulations were used to produce Figure 4 and Supplementary Figures 7, 8, 9, and 13.

To obtain a sense of the size of basins of attraction towards institutional adherence in finite populations, we repeated the previous set of simulations, except that, after reputation equilibration,

a random number $\in [1, N-1]$ of individuals was introduced with institution-adherence (type Z_a). A total of 2,500 replicate simulations was performed for each parameter combination and permitted to run until the adherence trait was driven to either fixation or extinction. We recorded the mean fixation frequency over these replicate simulations. These simulations were used to produce Supplementary Figures 1 and 2.

Finally, we conducted simulations to estimate fixation probability of novel mutant adherent (Z_a) introduced into a population containing a mixture of non-adherent discriminators (Z_n), unconditional cooperators (X), and unconditional defectors (Y). The protocol for these simulations was identical to that described in the paragraph above, except that ALLC and ALLD players do not condition their behavior on reputations. We estimated fixation probability starting from four different configurations, corresponding to different positions on the X, Y, Z_a simplex: simulations in which the initial frequencies were set to 16/50, 17/50, and 17/50 respectively (roughly the center of the simplex), and three sets of simulations, each starting with one strategy frequency set to 46/50 and the other two set to 2/50 (near the corners of the simplex). Reputations were permitted to equilibrate for 100 generations without strategy updating. A single DISC-PRIVATE individual Z_n was then switched to become DISC-ADHERE, and the simulation was permitted to run until DISC-ADHERE either fixed or went extinct. A total of 2,500 replicate simulations were used to estimate the fixation probability of DISC-ADHERE under these conditions. These simulations were used to produce Supplementary Figures 3, 4, 5, and 6.

Data and Code Availability

There are no empirical data associated with this study. All computer code for the simulations that produce figures are deposited as part of the publication.

Supplementary Figures

Supplementary Figure 1. Evolution can stabilize public monitoring institutions under all social norms. We plot the mean fixation probability of a mutant type that follows the institution (DISC-ADHERE) in a population that otherwise uses private assessment (DISC-PRIVATE), averaged over an initial mutant frequency f_{Z_a} drawn uniformly from the open interval $(0, 1)$. Vertical dotted lines represent mean frequencies of “good” reputations in the population at time of mutant introduction, averaged over all subjective perspectives. Horizontal lines (0.5) indicate the mean fixation probability of a neutral mutant. The institution is comprised of an external board.

Supplementary Figure 2. Evolution can stabilize public monitoring institutions under all social norms. All parameter values and notations are identical to Supplementary Figure 1, except that in this figure the institution is comprised of an internal board.

Supplementary Figure 3. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population three non-adherent types (ALLC, X ; ALLD, Y ; and DISC-PRIVATE, Z_n) initially each at frequencies 16/50, 17/50, and 17/50, respectively. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 4. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population of three non-adherent types (ALLC, X ; ALLD, Y ; and DISC-PRIVATE, Z_n) with initial frequencies $f_X = 2/50$, $f_Y = 2/50$, and $f_{Z_n} = 46/50$. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 5. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population of three non-adherent types ((ALLC, X ; ALLD, Y ; and DISC-PRIVATE, Z_n) with initial frequencies $f_X = 2/50$, $f_Y = 46/50$, and $f_{Z_n} = 2/50$. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 6. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population of three non-adherent types (ALLC, X ; ALLD, Y ; and DISC-PRIVATE, Z_n) with initial frequencies $f_X = 46/50$, $f_Y = 2/50$, and $f_{Z_n} = 2/50$. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 7. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population of non-adherents (DISC-PRIVATE, Z_n), with $b = 5$ and an internal board. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 8. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population of non-adherents (DISC-PRIVATE, Z_n), with $b = 60$ and an external board. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 9. Fixation probability for a single institutional adherent (DISC-ADHERE, Z_a) introduced in a population of non-adherents (DISC-PRIVATE, Z_n), with $b = 60$ and an internal board. Parameter values are otherwise the same as in Figure 4.

Supplementary Figure 10. Selection gradient of strategy evolution under replicator dynamics. All notation and parameter values are identical to Figure 2, except that in this case e_1 exceeds e_2 : $e_1 = 0.1$ and $e_2 = 0.01$.

Supplementary Figure 11. Selection gradient of strategy evolution under replicator dynamics. All notation and parameter values are identical to Figure 2, except that in this case e_2 exceeds e_1 : $e_1 = 0.01$ and $e_2 = 0.1$.

Supplementary Figure 12. Selection gradient of strategy evolution under replicator dynamics. All notation and parameter values are identical to Figure 2, except that in this the benefit of cooperation is larger: $b = 10$.

Supplementary Figure 13. Growth rate of institutional adherents (DISC-ADHERE, Z_a) in a population otherwise composed of private assessors (DISC-PRIVATE, Z_n). All parameters and notation are the same as in Figure 5, except that the institution here consists of an internal board.

References

- Richard D Alexander. A biological interpretation of moral systems. *Zygon*®, 20(1):3–20, 1985.
- Tamas Bereczkei, Bela Birkas, and Zsuzsanna Kerekes. Public charity offer as a proximate factor of evolved reputation-building strategy: an experimental analysis of a real-life situation. *Evolution and Human Behavior*, 28(4):277–284, 2007.
- Oliver Scott Curry, Daniel Austin Mullins, and Harvey Whitehouse. Is it good to cooperate. *Curr Anthropol*, 60(1):47–69, 2019.
- Jon Elster. Social norms and economic theory. *Journal of economic perspectives*, 3(4):99–117, 1989.
- Ernst Fehr and Urs Fischbacher. Social norms and human cooperation. *Trends in cognitive sciences*, 8(4):185–190, 2004.
- Ernst Fehr, Urs Fischbacher, and Simon Gächter. Strong reciprocity, human cooperation, and the enforcement of social norms. *Human nature*, 13(1):1–25, 2002.
- Bruno S Frey. How intrinsic motivation is crowded out and in. *Rationality and society*, 6(3):334–352, 1994.
- Bruno S Frey and Reto Jegen. Motivation crowding theory. *Journal of economic surveys*, 15(5):589–611, 2001.
- Bruno S Frey and Felix Oberholzer-Gee. The cost of price incentives: An empirical analysis of motivation crowding-out. *The American economic review*, 87(4):746–755, 1997.
- Uta Frith and Francesca Happé. Theory of mind and self-consciousness: What is it like to be autistic? *Mind & language*, 14(1):82–89, 1999.
- Avner Greif. Impersonal exchange without impartial law: the community responsibility system. *Chi. J. Int’l L.*, 5:109, 2004.
- Michael Hechter. Norms in the evolution of social order. *Social Research: An International Quarterly*, 85(1):23–51, 2018.
- Christian Hilbe, Laura Schmid, Josef Tkadlec, Krishnendu Chatterjee, and Martin A Nowak. Indirect reciprocity with private, noisy, and incomplete information. *Proceedings of the National Academy of Sciences*, 115(48):12241–12246, 2018.
- Michihiro Kandori. Social norms and community enforcement. *The Review of Economic Studies*, 59(1):63–80, 1992.
- Justin W Martin, Jillian J Jordan, David G Rand, and Fiery Cushman. When do we punish people who don’t? *Cognition*, 193:104040, 2019.
- Naoki Masuda and Francisco C Santos. Social psychology: A mathematical look at empathy. *eLife*, 8:e47036, 2019.
- Craig McIntosh, Elisabeth Sadoulet, Steven Buck, and Tomas Rosada. Reputation in a public goods game: Taking the design of credit bureaus to the lab. *Journal of Economic Behavior & Organization*, 95:270–285, 2013.

- Ruth Meinzen-Dick. Beyond panaceas in water institutions. *Proceedings of the national Academy of sciences*, 104(39):15200–15205, 2007.
- Paul R Milgrom, Douglass C North, and Barry R Weingast*. The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs. *Economics & Politics*, 2(1):1–23, 1990.
- Mitsuhiko Nakamura and Naoki Masuda. Indirect reciprocity under incomplete observation. *PLoS computational biology*, 7(7):e1002113, 2011.
- Martin A Nowak and Karl Sigmund. Evolution of indirect reciprocity. *Nature*, 437(7063):1291, 2005.
- Hisashi Ohtsuki and Yoh Iwasa. How should we define goodness?—reputation dynamics in indirect reciprocity. *Journal of theoretical biology*, 231(1):107–120, 2004.
- Hisashi Ohtsuki and Yoh Iwasa. The leading eight: social norms that can maintain cooperation by indirect reciprocity. *Journal of Theoretical Biology*, 239(4):435–444, 2006.
- Hisashi Ohtsuki, Yoh Iwasa, and Martin A Nowak. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature*, 457(7225):79, 2009.
- Isamu Okada, Tatsuya Sasaki, and Yutaka Nakai. Tolerant indirect reciprocity can boost social welfare through solidarity with unconditional cooperators in private monitoring. *Scientific reports*, 7(1):9737, 2017.
- Isamu Okada, Tatsuya Sasaki, and Yutaka Nakai. A solution for private assessment in indirect reciprocity using solitary observation. *Journal of theoretical biology*, 455:7–15, 2018.
- Elinor Ostrom. Collective action and the evolution of social norms. *Journal of economic perspectives*, 14(3):137–158, 2000.
- Jorge M Pacheco, Francisco C Santos, and Fabio AC C Chalub. Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS computational biology*, 2(12):e178, 2006.
- Arunas L Radzvilavicius, Alexander J Stewart, and Joshua B Plotkin. Evolution of empathetic moral evaluation. *eLife*, 8:e44269, 2019.
- David G Rand, Corina E Tarnita, Hisashi Ohtsuki, and Martin A Nowak. Evolution of fairness in the one-shot anonymous ultimatum game. *Proceedings of the National Academy of Sciences*, 110(7):2581–2586, 2013.
- Mari Rege and Kjetil Telle. The impact of social approval and framing on cooperation in public good situations. *Journal of public Economics*, 88(7-8):1625–1644, 2004.
- Devesh Rustagi, Stefanie Engel, and Michael Kosfeld. Conditional cooperation and costly monitoring explain success in forest commons management. *science*, 330(6006):961–965, 2010.
- Fernando P Santos, Jorge M Pacheco, and Francisco C Santos. Evolution of cooperation under indirect reciprocity and arbitrary exploration rates. *Scientific reports*, 6:37517, 2016.
- Fernando P Santos, Francisco C Santos, and Jorge M Pacheco. Social norm complexity and past reputations in the evolution of cooperation. *Nature*, 555(7695):242, 2018.

- Tatsuya Sasaki, Isamu Okada, and Yutaka Nakai. The evolution of conditional moral assessment in indirect reciprocity. *Scientific reports*, 7:41870, 2017.
- Karl Sigmund. Moral assessment in indirect reciprocity. *Journal of theoretical biology*, 299:25–30, 2012.
- Adam Smith. Cognitive empathy and emotional empathy in human behavior and evolution. *The Psychological Record*, 56(1):3–21, 2006.
- Kristopher M Smith and Coren L Apicella. Hadza hunter-gatherers disagree on perceptions of moral character. *Social Psychological and Personality Science*, page 1948550619865051, 2019.
- Ralf D Sommerfeld, Hans-Jürgen Krambeck, Dirk Semmann, and Manfred Milinski. Gossip as an alternative for direct observation in games of indirect reciprocity. *Proceedings of the national academy of sciences*, 104(44):17435–17440, 2007.
- Peter D Taylor and Leo B Jonker. Evolutionary stable strategies and game dynamics. *Mathematical biosciences*, 40(1-2):145–156, 1978.
- Michael Tomasello and Amrisha Vaish. Origins of human cooperation and morality. *Annual review of psychology*, 64:231–255, 2013.
- Arne Traulsen, Jorge M Pacheco, and Martin A Nowak. Pairwise comparison and selection temperature in evolutionary game dynamics. *Journal of theoretical biology*, 246(3):522–529, 2007.
- Arne Traulsen, Dirk Semmann, Ralf D Sommerfeld, Hans-Jürgen Krambeck, and Manfred Milinski. Human strategy updating in evolutionary games. *Proceedings of the National Academy of Sciences*, 107(7):2962–2966, 2010.
- Satoshi Uchida and Karl Sigmund. The competition of assessment rules for indirect reciprocity. *Journal of Theoretical Biology*, 263(1):13–19, 2010.
- Sonja Utz, Uwe Matzat, and Chris Snijders. On-line reputation systems: The effects of feedback comments and reactions on building and rebuilding trust in on-line auctions. *International Journal of Electronic Commerce*, 13(3):95–118, 2009.
- Christopher R von Rueden, Daniel Redhead, Rick O’Gorman, Hillard Kaplan, and Michael Gurven. The dynamics of men’s cooperation and social status in a small-scale society. *Proceedings of the Royal Society B*, 286(1908):20191367, 2019.
- Hitoshi Yamamoto, Isamu Okada, Satoshi Uchida, and Tatsuya Sasaki. A norm knockout method on indirect reciprocity to reveal indispensable norms. *Scientific reports*, 7:44146, 2017.